

NIF ICCS Network Design and Loading Analysis

R. Bryant

U.S. Department of Energy

February 20, 1998

Lawrence
Livermore
National
Laboratory

DISCLAIMER

This document was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor the University of California nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or the University of California, and shall not be used for advertising or product endorsement purposes.

Work performed under the auspices of the U. S. Department of Energy by the University of California Lawrence Livermore National Laboratory under Contract W-7405-Eng-48.

This report has been reproduced
directly from the best available copy.

Available to DOE and DOE contractors from the
Office of Scientific and Technical Information
P.O. Box 62, Oak Ridge, TN 37831
Prices available from (423) 576-8401
<http://apollo.osti.gov/bridge/>

Available to the public from the
National Technical Information Service
U.S. Department of Commerce
5285 Port Royal Rd.,
Springfield, VA 22161
<http://www.ntis.gov/>

OR

Lawrence Livermore National Laboratory
Technical Information Department's Digital Library
<http://www.llnl.gov/tid/Library.html>

NIF ICCS Network Design and Loading Analysis

NIF-004126

Robert Bryant,
February 20, 1998

Introduction

The National Ignition Facility (NIF) is housed within a large facility about the size of two football fields. The Integrated Computer Control System (ICCS) is distributed throughout this facility and requires the integration of about 40,000 control points and over 500 video sources. This integration is provided by approximately 700 control computers distributed throughout the NIF facility and a network that provides the communication infrastructure.

A main control room houses a set of seven computer consoles providing operator access and control of the various distributed "front-end processors" (FEPs). There are also "remote workstations" distributed within the facility that allow provide operator console functions while personnel are testing and troubleshooting throughout the facility. The operator workstations communicate with the FEPs which implement the localized control and monitoring functions. There are different types of FEPs for the various subsystems being controlled.

This report describes the design of the NIF ICCS network and how it meets the traffic loads that will be expected and the requirements of the Sub-System Design Requirements (SSDR's).

This document supersedes the earlier reports entitled "Analysis of the National Ignition Facility Network," dated November 6, 1996 and "The National Ignition Facility Digital Video and Control Network," dated July 9, 1996.

For an overview of the ICCS, refer to the document "NIF Integrated Computer Controls System Description" (NIF-3738).

Network Requirements

There are two sources of network requirements for the NIF network that affect the overall network design, Sub-System Design Requirements (SSDR's) and Interface Control Documents (ICD's).

NIF SSDR's define general system requirements that must be met within specific Work Breakdown Structure (WBS) areas. There are two SSDR's that identify network-related requirements, SSDR 1.5.1 (Computer System) and SSDR 1.5.6 (Ancillary Systems). The network-related requirements from these two documents are extracted and discussed in a later section.

NIF ICD's define the interfaces between the different subsystems as defined by WBS's. The specific network interfaces/requirements are defined in a set of 27 ICD's between WBS 1.5.1.2 (Integrated Computer Control/ Computer System/ Operations Computer System) and the other WBS's having network needs. The following list shows the WBS's for which ICDs exist relating to the network.

WBS #	System Name	System Contact
1.3.1	Optical Pulse Generation (OPG)	Deadrick
1.3.3.3	Plasma Electrode Pockels Cell FEP	Bettenhausen, Lagin
1.3.4.5	Power Conditioning	McKay
1.3.5	Laser Auxiliary Systems (Amp Cooling)	Karlsen/Reed
1.4.2	Vacuum and Gas Control System	Karlsen/Reed
1.5.2.2	Alignment Controls Supervisory Software	Saroyan
1.5.2.3	Power Conditioning Supervisory Software	Carey, Bettenhausen
1.5.2.4	Laser Diagnostics Supervisory Software	Carey
1.5.2.5	Target Diagnostics Supervisory Software	Saroyan
1.5.2.6	OPG Supervisory Software	Gorvad
1.5.2.7	Shot Director Supervisory Software	Kettering
1.5.3	Integrated Timing System	Wiedwald
1.5.4	Safety Integrated System	Reed
1.5.5	Automatic Alignment	Miller
1.5.6.1	Video System	Van Atta
1.5.6.3	Facility Environment Monitor	Reed
1.7.1.4	Alignment Controls Systems	Miller Kamm
1.7.2.3.2	Precision Diagnostic Station	Demaret
1.7.2.3.4&5	Optics Inspection System	Kyker
1.7.2.3.6	Pulse Synchronization Processor	Kyker
1.7.2.4	Energy Diagnostic Processor	Demaret
1.7.2.5	Power Diagnostics Processor	Demaret
1.7.3	Wavefront Control FEP	Van Atta
1.8.1	Target Vacuum System	Reed/Weed
1.8.3.3	Target Diagnostics Data Acquisition	Dohoney
1.8.5	Environmental Protection Systems	Dohoney/Reed
1.8.6.1	Target Area Local Utility Services	Dohoney
1.8.7	FOA Thermal System	Reed
1.9.1.2	OAB Supervisory & Control	Sander

The ICD's specify the following information:

- Types and quantities of network attached devices (e.g. 10 or 100 Mbit/s Ethernet, ATM)
- Physical attachment type (e.g. fiber - SC or ST connector, AUI, RJ-45)
- Network software interface (e.g. TCP/IP, ATM API)
- Peak bandwidth needs
 - Peak bandwidth time interval (during the shot interval)
 - To-from message rates and sizes
- Boot image size

- Initialization message rates and sizes
- Shot archive data size
- Shot history data size

The ICD information is detailed in a set of spreadsheets from which graphs are generated showing aggregate traffic rates during different time on the shot cycle. These ICD analysis spreadsheets are attached to this report. This information has been analyzed in several ways to ensure that the proposed network design meets the ICD requirements. A network modeling effort is also in process (funded by LDRD) which will allow us to further analyze the network design and perform "what-if" scenarios as new requirements are proposed and new hardware/software becomes available during NIF development and operation.

Network Traffic Types

There are three general types of traffic that will be supported by the network - data, video, and triggers.

The data traffic includes the transfer of information such as commands and data that do not require tight bounds on network latency (i.e. end-to-end delay) and jitter (i.e. end-to-end delay variation). This includes control messages, sensor data, boot images, archive and history data, etc.

The video traffic comes from the 500+ video cameras located throughout the NIF facility which are front-ended by 28 Video FEP's. A Video FEP can interface to 24 cameras and can selectively grab frames from these cameras on demand from the operator/remote consoles or the automatic alignment system. The FEP can send one frame on request or send a stream of frames at up to 10 frames/sec for real time monitoring. All digitized frames are uncompressed 640 x 480 x 8 bits or about 2.5 Mbits per frame. Compression is not being utilized because of the additional processing burden and cost it would impose on the sending and receiving systems. At 10 frames/sec, a video stream requires approximately 25 Mbits/s of bandwidth. The video stream will be multicast when multiple operator/remote consoles want to view the same camera. Each operator/remote workstation will be capable of displaying one video stream into a window on each of its attached monitors. The main console workstations have two monitors each and the remote workstations have a single monitor. Each video FEP will be able to source up to 3 video streams simultaneously.

The network trigger traffic consists of short messages requiring a short "guaranteed" network latency (e.g. less than 10 msec) which are multicast to a specific set of end systems. The recipient end systems use these triggers to initiate a time critical function, such as just prior to shot time for triggering the video FEPs to grab the next video and the T-1 trigger to prepare the alignment control systems for the imminent shot.

The large majority of the traffic flow is between a few centralized systems (i.e. the supervisory consoles, file servers, automatic alignment systems, and the

master timing system) and the hundreds of other "distributed" systems located throughout the NIF facility. There is little traffic between FEPs.

Network Booting, Initialization, Archive, and History

Prior to the approximately seven hour shot cycle, all end systems must be bootstrapped and initialized. Since most of the end systems will be diskless, it is necessary that they obtain their boot images from the centralized file servers. The requirement is to boot and initialize all systems with 30 minutes. The time to boot is strongly dependent on the approach taken for booting. A hierarchical approach where the file servers boot other servers which then boot their FEPs will require much less time than if the file servers send boot images directly to all systems. For a discussion of the boot process and analysis, see "Startup Simulation," NIF #004132.

After the shot has occurred, archive and history data that was collected by the various FEP's needs to be saved to the file servers. The goal is to transfer all the archive and history data within 30 minutes. The network has been designed to provide sufficient bandwidth to meet this requirement.

Equipment Locations

The centralized computer systems (i.e. the supervisory consoles, file servers, automatic alignment systems, and the master timing system) are located in the Main Control room, Computer room, Classified Control room, and the Master Oscillator room (MOR). The distributed systems are located in the Laser Bays, Switchyards, Mezzanines, Target Bay, Diagnostic building, Capacitor Bays, MOR, PAMMA room, and Optics Assembly building.

Specific rack or enclosure information was collected for each of the network attached devices throughout NIF. Based on this information, locations for distributed Ethernet switches and fiber patch panels were selected to provide network services within a localized area.

Network Design

The network design is summarized in the attached NIF Network Diagram. More detailed schematics are provided in the set of drawings LEA96-2821.

The design utilizes both Ethernet (10 and 100 Mbit/s) and Asynchronous Transfer Mode (ATM) technologies to take advantage of the best features of each. ATM is utilized on systems having time-sensitive video requirements which includes the video FEPs, the operator and remote workstations, and the automatic alignment, master timing, target diagnostics, and optics inspection systems. Ethernet provides connectivity to all the systems in the network even if they also have ATM connections. Since the centralized computers need to communicate with ATM and Ethernet attached FEPs, the operator consoles and file servers have interfaces to both ATM and Ethernet (i.e. they are dual-homed). Many of the FEPs are diskless, so Ethernet connections are required on those systems to allow booting (i.e. an ATM bootstrap ROM does not exist).

The following table lists some key advantages and disadvantages of Ethernet and ATM technologies.

	Ethernet	ATM
NIF Usage	For all traffic except video-related	For video-related traffic
Advantages	Very low cost	QoS control over video streams reduces jitter and improves stream delivery (can reserve bandwidth)
	10 Mb/s is standard on all systems; 100 Mb/s is standard on most new systems	Can use larger packet sizes for direct ATM transfers (up to 64 KB) for greater efficiency
	Simple to work with (hardware and software configuration)	
	Efficient for small message traffic	
	Gbit Ethernet coming	
Disadvantages	Current systems do not support Quality of Service (QoS) which is useful for video streams	Much more complex to configure (lots of options)
	Doesn't support large packet sizes (1,500 Byte max.), requiring more packet processing overhead	Requires another card and device driver to be installed adding cost and complexity

Core high-capacity 100 Mbit/s Fast Ethernet and 155 Mbit/s OC-3 ATM switching systems are located in the computer room. The Core Fast Ethernet switch connects to the centralized servers and control consoles, and to the distributed 10 and 100 Mbit/s switches located throughout the NIF facility. The centralized ATM switch is connected via fiber directly to the end systems throughout NIF except in the Target Diagnostics area where a small ATM switch provides multiplexing.

A following table lists port usage on the Core ATM switch, Core Ethernet switch, and distributed Ethernet switches.

Switch type	Port type	Qty	System
-------------	-----------	-----	--------

Core ATM Switch	155 Mb/s	62	14	Operator Workstations
	OC-3			2 File Servers
				14 Remote Workstations
				24 Video FEPs
				3 Automatic Alignment systems
				1 Master Timing FEP
				1 Optics Inspection FODI FEP
				1 Target Diagnostics
				1 Classified Control rm console
				1 Core Ethernet Switch
Core Ethernet Switch	100 Mb/s	75	47	Distributed Ethernet Switches
				16 Operator Consoles
				2 File Servers
				3 Automatic Alignment systems
				1 Master Timing FEP
				1 Target Diagnostics
				1 Classified Control rm console
				1 Optics Assembly bldg
				1 Network Management station
				1 Firewall
Distributed Ethernet Switches	100 Mb/s	78	72	Alignment Control T-1 FEPs
				4 Power Conditioning FEPs
				2 Optics Inspection LODI FEPs
	10 Mb/s	57	48	Wavefront FEPs
		8	12	PEPC FEPs
			36	Alignment Controls FEPs
			64	Optical Pulse Gen. FEPs
			27	Zone Timing FEPs
			28	Video FEPs
			20	Power Diagnostics systems
			6	Energy Diagnostics FEPs
			22	Precision Diagnostics Station
			1	Optics Inspection FODI FEP
			44	Industrial Controls Systems
			19	Power Conditioning Controllers
			38	Optics Assembly Bldg
			40	Target Diagnostics

During the original Title I design phase, the network design had a pair of ATM switches at the core of the network through which all traffic flowed between the supervisory workstations, file servers, and automatic alignment system, and the Ethernet edge switches that connect to the distributed FEPs. Since then, Fast Ethernet products have become widespread and prices have dropped significantly, making it the best candidate for non-video data transport. A separate core Fast Ethernet switch now provides the

connectivity between the Ethernet edge switches and the centralized computer systems. There are many advantages of this approach, including:

- 1) It is a simpler, less complex architecture since Ethernet is used to communicate end-to-end. A translation between Ethernet and ATM using LAN Emulation (LANE) will not be required. Although LANE products have become more production-worthy in the last year, LANE requires much more configuration, management, and maintenance expertise to operate them. This would likely result in more operational problems. Ethernet is relatively simple and well understood.
- 2) Separating the Ethernet and ATM traffic eliminates the potential for interference between the latency-sensitive traffic (i.e. video) on the ATM network and the other traffic. The total bandwidth to the central computers is also essentially doubled in the process since they are dual-homed.
- 3) The performance of LANE tops out at just over 100 Mbit/s (based on our tests) and there is a significantly higher CPU utilization than Ethernet due to the extra complexity in the LANE client software.
- 4) Fast Ethernet ports are standard on most new products, whereas ATM must be purchased separately. All the central computers have Fast Ethernet ports and even 10 Mbit/s Ethernet switches come with Fast Ethernet uplinks at little extra cost. The cost of the distributed switches are significantly less with Fast Ethernet uplinks than with ATM uplinks.
- 5) The overall cost is reduced by using a Core Ethernet switch rather than using an ATM switch in its place.
- 6) Ethernet is gaining momentum while the niche for ATM as a LAN technology is being questioned. With Gbit Ethernet in development and Ethernet prices dropping very rapidly, it is clear that Ethernet will be widely deployed well into the future.

If bandwidth requirements increase in the future, the ICCS network architecture will allow the integration of Gbit/s Ethernet and 622 Mb/s ATM technologies in a relatively straightforward manner. There is independence between the Ethernet and the ATM portions of the network, allowing the movement of systems between them as the technologies evolve in capabilities and price.

Ethernet Network

The Ethernet network utilizes switching technology throughout. Switches have significant performance advantages over shared Ethernet hubs. Also, prices for high capacity switches have been dropping very rapidly and now are not substantially more expensive than shared hubs.

The core Fast Ethernet switch is a modular rack mounted chassis with both multimode fiber and twisted pair ports. It has a redundant power supply to provide added reliability. The fiber ports have SC type connectors and the twisted pair ports have RJ-45 connectors, which are standards in the industry. The switch also has an OC-3 ATM port supporting LANE which is connected to the core ATM switch. This switch will also have the capability to support Gbit Ethernet.

The distributed Ethernet switches will be small "stackable" units supporting between 12 and 24 ports of 10 and/or 100 Mbit/s each and a Fast Ethernet fiber uplink (100Base-FX) to the core Ethernet switch. Both twisted pair and fiber switches are used in the design. The twisted pair ports use RJ-45 connectors. The 10 Mbit/s fiber ports use ST connectors, which are the industry standard.

Unlike shared Ethernet hubs, through which only one attached system can communicate at a time, Ethernet switches allow multiple simultaneous communications by isolating traffic to only the ports involved. An Ethernet switch learns where ports are located by monitoring source addresses as traffic enters the switch. Due to high speed ASIC technology, modern switches can run at wire speeds (even for 64 byte minimum sized frames), thus providing essentially full 10 Mb/s or 100 Mb/s rates through its ports (discounting the small overhead for interframe gaps, headers, etc). Aggregate capacities through 100 Mb/s switches vary between products, but rates in excess of 1 Gb/s are typical, with multiple Gb/s in higher capacity enterprise switches (such as will be used as the NIF Core Switch).

Besides the throughput advantages of switches, they also have significant latency advantages. Because collisions don't occur and the ports can run at wire speed, the latencies through these switches are low. Very large latencies on shared hubs are possible due to collisions under heavy loads and what is known as the "capture effect" where one port sending large amounts of data can block another port from sending. Although the minimum possible latency on a hub is lower than a switch due to the store and forward time in a switch, unless input port traffic exceeds the link rate of an output port, latencies will be low and relatively consistent.

The latency difference has been demonstrated with tests run over a shared Ethernet (with a user traffic load) and a switched Fast Ethernet. Using a simple program and an analyzer that measured the time to send and receive messages, the latencies for a shared 10 Mb/s Ethernet varied from 1.65 msecs up to 102 msecs, whereas the switched 100 Mb/s connection varied only between 83 and 137 usecs. These times include both the time on the network and the protocol stack times in both the sender and receiver (which were running VXworks).

Because of the very low latencies exhibited by 100 Mb/s switches, it is possible to support that transport of triggers across the network. This will be discussed in a later section.

ATM Network

The ATM design utilizes a high capacity core ATM switch with OC-3 ports connected directly to the end systems, except in the Target Diagnostics area, where a small ATM switch provides connectivity to end devices in that area. The core ATM switch will support a minimum of 64 ports and have an aggregate capacity of 10 Gb/s, providing non-blocking capacity for all ports. It has the capability to support OC-12 ports (622 Mbit/s) if needed in the future. All ports have SC connectors for use over multimode fiber. The switch supports various classes of service including CBR, VBR, and UBR, and provides LANE Services for the attached LANE clients.

Protocols

TCP/IP is used when guaranteed error-free end-to-end delivery is required. TCP provides retransmission of packets in the event that one is lost or received in error. The only traffic not using TCP/IP will be digitized video and network triggers.

The only Ethernet traffic that would not use TCP/IP are triggers which rely on IP multicasting. The feasibility of sending the T-1 trigger over Fast Ethernet is being considered. In this approach, the Master Timing system would use UDP to send a trigger message to a predefined multicast IP address. Upon receiving this message, the Ethernet switches would broadcast the message simultaneously onto all its ports. Thus, the Alignment Control T-1 systems will all receive the message after only two Fast Ethernet switch latency times. Testing, simulation, and analysis all indicate that the network latency will be under 1 msec under light to moderate network loads, which will be the case at the time the trigger is sent just before the shot.

Different protocols are used over the ATM network depending on the application.

Streaming digitized video sent between a video FEP and a supervisor console or remote workstation will use direct ATM API communication based on the XTI standard. This will maximize bandwidth usage, minimize CPU utilization on both ends, and allow for the establishment of point-to-multipoint switched virtual circuits (SVC's) with a guaranteed bandwidth for the video streams (25 Mb/s is needed for 640 x 480 x 8 x 10 frames/sec). All communication between the requesting workstation and the video FEPs for setting up and tearing down video streams will use TCP/IP because of its guaranteed end-to-end reliability. These control messages are small and are not time critical.

The video capture trigger will be sent over the ATM network from the Master Timing system to all the Video FEPs. This trigger will be sent using IP multicast over a point-to-multipoint PVC that is permanently configured in the ATM switch. Using this approach, the Master Timing system sends a single trigger message that is replicated in the switch and forwarded to all

the video FEPs nearly simultaneously. Tests have verified that the switch and NIC latencies are typically 150 microseconds, which provides a sufficient margin for delivery of the trigger within the specified 5 msecs.

All other traffic on the ATM network will use TCP/IP to provide guaranteed reliable end-to-end delivery of data. This includes the transfer of single camera images between the video FEPs and the Automatic Alignment (AA) systems and between the Final Optics Damage Inspection system and the AA systems.

Cable Plant

A combination of standard multi-mode 62.5 micron fiber optic cables and Category 5 twisted pair cables will be utilized in the network. The use of twisted pair cabling will be limited to Ethernet connections with short lengths where optical isolation is not critical, such as within a group of racks and between the main control room, classified control room and the computer room. Fiber cabling is used for all ATM connections, for connection of distributed Ethernet switches to the core Ethernet switch, and for Ethernet connections where optical isolation is important, including the PEPC FEPs, the PAM FEPs, the Power Conditioning controllers, and the MOR network connections.

For Ethernet connections, there are some important advantages to using twisted pair cabling rather than fiber for short runs where optical isolation is not a requirement. The costs for implementing a fiber based Ethernet is approximately a factor of 7 more expensive than a twisted pair approach for both 10 and 100 Mb/s Ethernet. A cost comparison is shown at the end of this document.

Besides cost, there are other disadvantages to using fiber rather than Category 5 twisted pair.

- 1) 10Base-T and 100Base-TX to fiber optic media converters are needed since all systems come with 10/100Base-T ports as their base configuration. These converters require their own power supply which will take up another 110V outlet (or two depending on the size of the power "plug") and will create extra clutter in the racks. This clutter and extra cabling could lead to their own maintenance/reliability problems. Using modular rack mounted converters is neater, but also takes up extra rack space.

- 2) Fiber Ethernet switch ports can only support either 10 Mb/s or 100 Mb/s, not both as can twisted pair 10/100 Ethernet switches. Therefore, a fiber switch cannot have a mix of 10 and 100 Mbit/s ports, nor transparently allow end systems to upgrade to from 10 to 100 Mbit/s without upgrading both switches and fiber converters. This makes network upgrades more complex and expensive.

3) The density of Ethernet fiber hubs and switches are typically a factor of 2 lower than twisted pair hubs, requiring a bit more rack space and power.

4) The choices of vendors and products for Ethernet fiber hubs and switches are limited, which doesn't allow as much flexibility of choice and price competition.

Fiber patch panels will be used in the computer room and in the network equipment racks distributed throughout the facility. This will ensure the reliability and maintainability of the network and also keep the installation tidy. Spare fibers and patch points are provided allowing for the addition of new capabilities in the future. ST connectors will be used in these patch panels.

Analysis of Network Requirements from SSDR 1.5.1 & 1.5.6

This section describes how each of the network related requirements in SSDR 1.5.1 and 1.5.6 will be met by the NIF ICCS network and provides the supporting analysis.

Throughout this section, the NIF ICCS network will be referred to as "the network." The network includes the hardware, cabling, and software needed to communicate data between ICCS end systems. This includes the physical, link, network, and transport layers, which includes the TCP/IP protocol stack. In the discussion below about network performance, only the network protocol performance is discussed, not application performance which will also affect overall ICCS performance.

SSDR 1.5.6

Paragraph 3.2.03.1

Facility Wide Computer Network - Provide a facility-wide data network as a backbone for data communications in NIF. Several LANs shall be connected to this backbone. Subsystems with related control and data dependence shall be closely coupled on separately bridged and isolated segments to achieve desired performance and information transfer rates.

Refer to drawing package LEA96-2821 for detailed diagrams of the network design.

The NIF ICCS network interconnects all the control, diagnostic, server, and console systems as specified in the Interface Control Documents (ICDs) for WBS 1.5.1.2. The ICD's define the interfaces to the hundreds of systems that will be attached to the network. The network will interconnect systems throughout the NIF facility, including the laser bays, switchyards, upper and lower mezzanines, target area, control room, computer room, master oscillator room (MOR), PAMMA room, capacitor bays, diagnostic building, optics assembly building, and utility and BTS pump pads.

The core network backbone will have 100 Mbit/s Fast Ethernet and 155 Mbit/s Asynchronous Transfer Mode (ATM) switches. A core ATM switching system with 155 Mbit/s (OC-3c) ports will be used for digitized video traffic requiring low data transfer jitter and high bandwidth for multiple video streams. A core Fast Ethernet switching system will be used for all other traffic. It will connect to other Ethernet switches with 10 and 100 Mbit/s ports distributed throughout the facility providing localized connectivity to the end systems. End systems needing higher speed or lower latencies than can be provided by 10 Mbit/s Ethernet will be connected to 100 Mbit/s switch ports.

Ethernet switches will be used throughout the network, providing isolation of traffic between the switch ports. Ethernet switches filter/isolate traffic that is not destined for a device attached to a particular port.

The Ethernet switches will support virtual LAN (VLAN) capabilities allowing end systems that need to communicate with each other to be placed in separate virtual networks. This can be applied to further isolate traffic and reduce the extent of broadcast messages.

Paragraph 3.2.03.2

Computer Network, Safety Systems - No safety related controls shall operate over the facility-wide network.

See "NIF Safety Interlock Systems, Programmable Logic Controller Requirements," NIF #000525.

All direct safety controls are contained in the safety control system. The safety system uses the ICCS network to provide graphical user interfaces to the shot director and operations personnel.

Paragraph 3.2.04

Optical Communications Safety - Where the Computer System uses lasers for optical communication, it shall conform to applicable eye-hazard safety requirements to mitigate the possibility of eye exposure to an uncapped laser diode or unterminated active fiber source.

The fiber optic ports on OC-3 ATM and Fast Ethernet devices use LEDs rather than lasers. Although none are currently planned, any lasers that may be used in the network equipment will be Class 1 lasers designed for use in LANs over multimode fiber. Class 1 lasers are defined as products which do not permit human access to laser radiation in excess of the accessible limits of Class 1 for applicable wavelengths and durations. These lasers are safe under reasonably foreseeable conditions of operation. Also, personnel will adhere to LLNL Health & Safety Standards.

Paragraph 3.2.05

Computer System and Network Reliability - The network shall meet reliability standards for industrial control systems. Redundancy is not considered essential to achieve adequate reliability.

See Title II Design Basis Book - Section 13 FMEA Analysis.

The equipment and cabling used in the network are all commercial off-the-shelf and are extremely reliable systems. MTBF's are measured in years or 10's of years. A list of MTBF's for various network switches is shown at the end of this document.

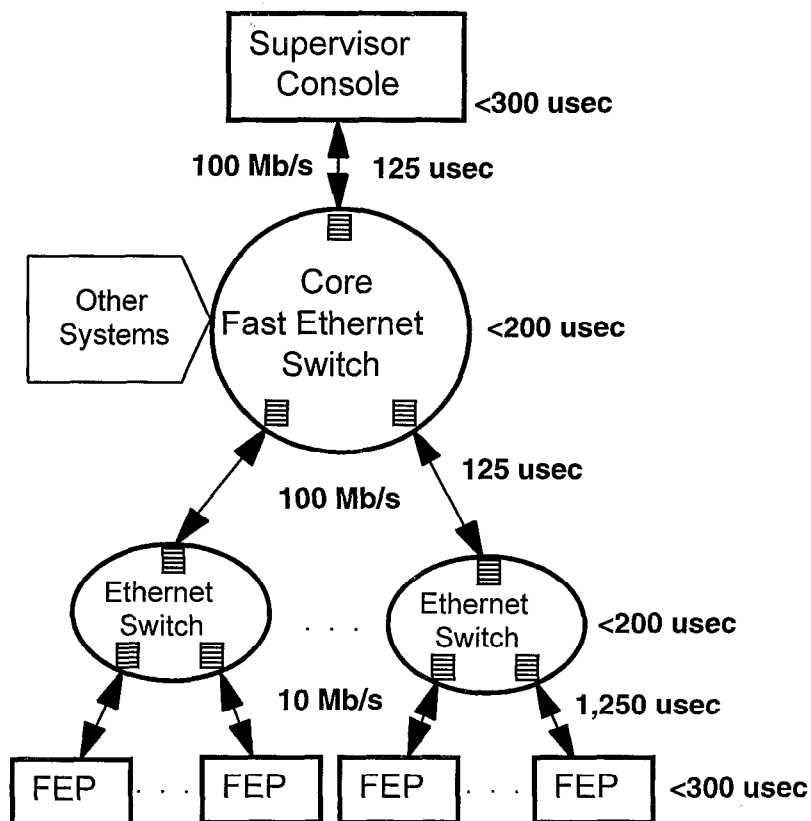
Redundancy will only be applied where the benefits outweigh the cost. The core network switches are the most critical so redundant power supplies will be provided.

Paragraph 3.2.07

Network Performance, Average Latency - The network's average two-way (round trip) time shall be less than 50 msec.

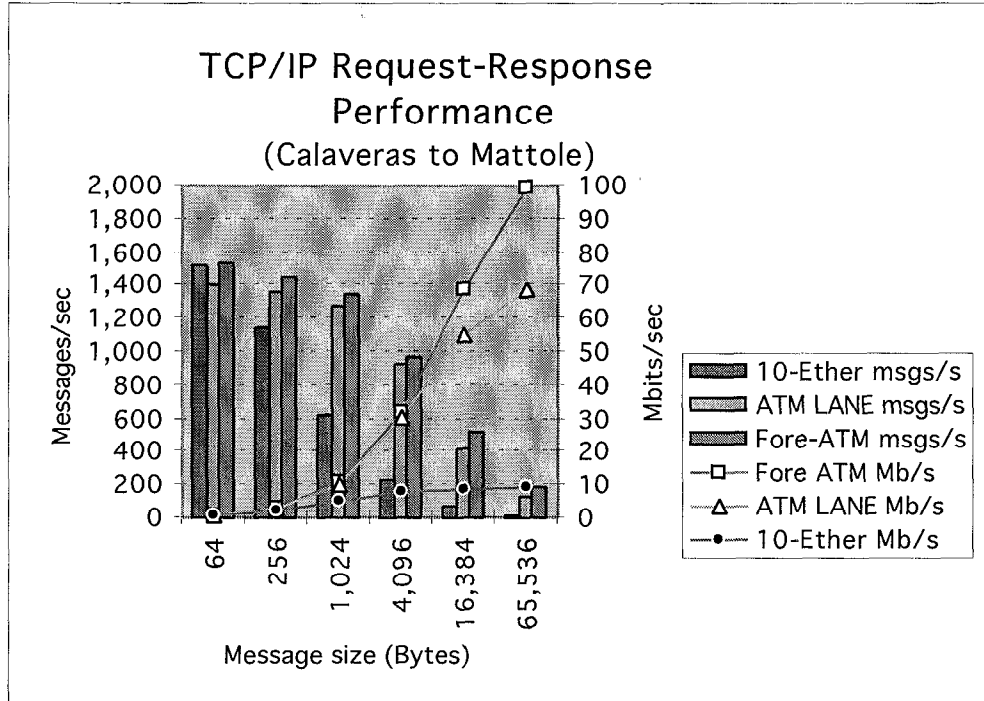
This requirement is easily met through the application of switching technology (rather than shared media devices) and an architecture that in most cases has only two Ethernet switches between communicating end systems. Switches isolate traffic between its ports such that performance between a pair of systems is not impacted by traffic between other systems attached to the switch. Current 10 Mb/s and even 100 Mb/s Ethernet switches have enough internal bandwidth to maintain "wire-speed" performance between all their ports. Based on various performance tests, simulations, and other analysis, the average two-way network latencies will easily be less than 5 msec. As mentioned earlier, this does not include latency within the applications running on the end systems.

The network round trip time can be divided into the network protocol processing in the end stations and the time traversing the network (switches and cabling). It is the sum of the one-way end-to-end latencies in each direction. The following figure diagrams the typical end-to-end configuration and latencies attributed to the various portions of the network for a maximum sized Ethernet packet of 1,500 bytes (which is longer than the majority of messages that will be sent). The latencies listed are conservative and are based on the relatively light network loading indicated by the ICD's.



Totaling these latencies, the total round trip latency will easily be less than 5 msec.

Results of actual performance tests using the Netperf Request-Response test between a pair of workstations in our Testbed supports this and are shown in the graph below. Over the 10 Mb/s Ethernet using TCP/IP with 1,024 byte messages, we were able to sustain more than 600 messages/sec. In this test, messages of the size shown were sent and a 1 byte message response was returned before the next message was initiated. These results indicate an average round trip latency of less than 2 msec over a shared 10 Mb/s Ethernet. With two layers of switching in the network as shown in the above figure, less than 2 msec of latency will be added, indicating a total latency of less than 4 msec.



The above analysis and tests were done on a lightly loaded network. What effect will loading have on the average latency? Our analysis indicates that network loading will have only a small effect on the round trip latency. This is because Ethernet switches isolate traffic between ports, so only traffic going to the same destination will affect the latency. Since the core Ethernet switch has 100 Mb/s ports which will be lightly loaded, there will be little increase in latency over what is described above.

In any case, the average round trip network latency will be \ll 50 msecs.

Paragraph 3.2.08

Network Performance, Maximum Response Time - The maximum network two-way (round trip) time maximum response time of less than 500 msecs under maximum loading

This requirement is also easily met through the use of switching technology and the 100 Mb/s core Ethernet switch. This keeps traffic isolated to the links that are in the path between the end systems. Where the peak traffic on all links is less than the available link bandwidth, then packets do not queue-up in output port buffers and the latency through a switch network is approximately the average rate, which was discussed in the previous section. There will be only a small variation through a switch even if a few packets destined for the same destination arrive at the same time.

Based on ICD traffic information, the peak traffic rates on all links during a shot cycle (-7 hr to shot time) are less than 10% of their rated capacity. In

most cases, the links will be running at less than 5% of capacity even at peak times. Therefore, it will be infrequent that messages will be delayed in the switch buffer queues due to other traffic going to the same destination.

In order for the round trip time to exceed 500 msecs, the one-way latency would have to exceed 250 msecs. That is, the sum of the protocol stack delay at the sender, the delay through the network, and the receiver stack delay would have to be greater than 250 msecs. The Ethernet packet latency across a link is 1.2 msec at 10 Mb/s and 120 microseconds at 100 Mb/s. To obtain a 250 msec. latency would require that a message be delayed in output queues behind over 200 maximum 10 Mb/s Ethernet sized packets in its path. Since the traffic on any given link will be much less than it's rated speed, there is no way that such queues could develop within any of the switches.

In our simulation tests, we were not able to cause latencies to grow significantly unless we increased the traffic levels from several sources to a particular destination port such that the destination link/port was overloaded. To achieve this result required traffic levels well beyond what will exist and what the network is designed to handle.

Paragraph 3.2.09

Network Performance, Sustained Data Rate - The network shall be capable of sustained throughput of greater than 100 KBytes/sec to any single end point. The use of subnets, filtering bridges and/or switches is suggested to maintain these rates.

Each network link is minimally a switched 10 Mbit/s Ethernet link which is capable of greater than 1 MByte/sec sustained throughput, based on our Netperf tests and vendor product information. With switching technology throughout the network and a 100 Mbit/s core switch, the network has sufficient capacity to sustain 100 KBytes/sec between any pair of systems.

Paragraph 3.2.10

Ethernet Loading - Where shared Ethernet is utilized, the average load shall be no greater than 15% of maximum over a shot cycle.

The network will utilize Ethernet switches throughout, with no shared Ethernet hubs planned. Even on the switched Ethernet links, utilization will be less than 10% overall.

Paragraph 3.2.20

Console System Flexibility - Any system shall be capable of displaying and operating any control subsystem.

The network is fully interconnected (i.e. every end system is capable of communicating with every other system) and is symmetric in that all console systems have identical network connections. Therefore, from a network

perspective, any console system is capable of displaying and operating any control subsystem.

Paragraph 3.2.21

Network Backbone - A low latency, high data rate (at least 100 Mb/s) network shall make up the digital communication system backbone.

The backbone utilizes 100 Mbit/s Ethernet and 155 Mbit/s ATM switches with aggregate capacities of greater than 3 Gbit/s and 10 Gbit/s respectively.

The physical cabling shall be configured in a hub and spoke (radial) fashion.

All cabling radiates from the network switches. Cabling runs out from the core switches and patch panels in the computer room to the other areas within NIF. Where distributed switches are used, cabling runs from them to the attached devices.

Paragraph 3.2.22

LAN Subnets (Ethernet) - The computer system shall provide LAN subnets using 10 Mbit/s Ethernet, if latency and throughput is adequate...10 Mbit/s Ethernet as the lowest common denominator network interface.

10 Mbit/s Ethernet is the slowest network interface used. Where higher performance is needed for improved performance or latency, 100 Mbit/s Ethernet or 155 Mbit/s ATM is utilized.

Paragraph 3.2.23

Lower-Level Computer Attachment to LAN - All equipment shall utilize network hardware and protocols compatible with the NIF network and bandwidth appropriate to the equipment requirements.

Interface Control Documents (ICD's) define the network hardware and protocols to be used as well as the bandwidth needs of the equipment. The network design is based on these ICD's ensuring that the network meets interface and performance needs.

Paragraph 3.2.24

Ethernet Protocol - At a minimum, the network shall support TCP/IP as the network transport protocol.

All network links support the TCP/IP protocol. There are also other protocols that will be used for specific applications, including UDP/IP (for multicasting/broadcasting), and direct ATM API communication (for streaming video)

Paragraph 3.2.28

Computer Availability, cold start - Restarting all computers from a power-down condition shall take less than 30 minutes.

See "Startup Simulation" NIF #004132 for simulation results of the startup process. The network has adequate bandwidth capacity to handle all the traffic which will occur during bootstrapping and initialization. This is based upon the analysis of boot image sizes and initialization message information provided in the ICDs. The file servers and supervisory workstations all have 100 Mbit/s Ethernet and 155 Mbit/s ATM interfaces connecting directly to the core switches. This allows the file servers to boot the supervisory workstations, which can then simultaneously boot the FEP's.

SSDR 1.5.6

Paragraph 3.1.1.1

System Description, TV Distribution System - The television distribution system routes approximately 1,000 alignment TV sensors to automatic alignment system FEP's and to operator console monitors where manual alignment is performed.

The network design allows for the planned 28 Video FEPs (VFEPs) that can connect up to 24 cameras each, for up to 672 cameras total. Each VFEP has a 155 Mbit/s OC-3 ATM link to the core ATM switch which has OC-3 links to each of the automatic alignment systems and each operator workstation and remote workstation. Through the ATM switch, video data from any video FEP can be routed to any of the automatic alignment systems and operator workstations upon request from an automatic alignment system, an operator workstation, or a remote workstation.

A preemptive routing architecture and video data compression reduces the total bandwidth required without imposing significant constraints.

The VFEPs will send single video images to the automatic alignment systems upon request and video streams to the operator workstations upon request. Data compression will not be used because of the additional processing that would be required on both ends to support it. The ATM network has adequate bandwidth to meet system requirements

Distribution of digitized video is via LAN

Digitized video is transported over the NIF ATM and Ethernet networks

Paragraph 3.2.00

Functional Requirements - The Television distribution system shall:

- provide routing capability to connect any input to any output
- provide TV monitoring in control consoles and at remote operator stations
- provide remote control of routing from computer via the LAN

See responses above for paragraph 3.1.1. In addition, all remote operator workstations are equipped with OC-3 ATM connections to the core ATM switch to support streamed video from any VFEP, similar to the supervisory control consoles.

Paragraph 3.2.08

Video Distribution - Ancillary System shall implement a video distribution system to route any video input to any monitor location.

See responses to paragraphs 3.1.1 and 3.2.00.

Motion Video Minimum Requirements:

- 2 operators per control room console (4 monitors) each viewing up to 2 images

Each operator console has two workstations. Each workstation has 2 monitors. Thus, there are 4 monitors per console. Each workstation has a 155 Mbit/s OC-3 ATM NIC over which the video streams are received. Each video stream of 640 x 480 x 8 bits at 10 frames/sec will utilize 25 Mb/s of network bandwidth. Two streams must be supported to each workstation to meet the requirement. This will utilize 50 Mbit/s of bandwidth on the 155 Mbit/s ATM link which is easily within the available bandwidth. Using the Netperf test code, we have demonstrated up to 54 video frames/sec are possible through the ATM API. Only 20 frames are required to meet the requirement.

- 10 frames per second at 480 by 512 pixels by 8 bits

The network will support the currently planned higher resolution of 640 by 480 pixels by 8 bits. As stated above, this will utilize 25 Mbit/s of bandwidth over an ATM link, which is easily supported and has been demonstrated.

- 2 cameras per beam must be viewable from one operator location at one time

Meeting this requirement is dependent upon factors outside the network, but the network will support the transport (multicasting) of the video streams for up to 3 cameras on any given video FEP. This will utilize 75 Mbit/s (more than half) of the available ATM link bandwidth. As mentioned above, each operator workstation will support the display of video streams from any 2 cameras.

Minimum number of monitor channels shall be 32

The network will support the transport of up to 3 video streams from each of the Video FEP's and the reception of 2 video streams on each of the operator workstations (in the main control room) and 1 video stream on any remote workstation. Since there are 28 video FEPs, 14 operator workstations and 14

remote workstations, the network is capable of supporting over 32 simultaneous video streams.

Paragraph 3.2.14

Optical Communications Safety - Where the Computer System uses lasers for optical communication, it shall conform to applicable eye-hazard safety requirements to mitigate the possibility of eye exposure to an uncapped laser diode or unterminated active fiber source.

The fiber optic ports on OC-3 ATM and Fast Ethernet devices use LEDs rather than lasers. Although none are currently planned, any lasers that may be used in the network equipment will be Class 1 lasers designed for use in LANs over multimode fiber. Class 1 lasers are defined as products which do not permit human access to laser radiation in excess of the accessible limits of Class 1 for applicable wavelengths and durations. These lasers are safe under reasonably foreseeable conditions of operation. Also, personnel will also adhere to LLNL Health & Safety Standards.

Twisted pair vs. fiber cabling cost comparison

<u>10 Mb/s cost/port</u>	<u>Twisted pair</u>	<u>Fiber</u>
Ethernet Hubs	\$50	\$220
End system cost	\$0	\$240 (10Base-T to fiber converter)
Cable cost (50 ft. cable)	<u>\$25 (Cat 5)</u>	<u>\$60</u>
Total	\$75	\$520

Assuming 500 10 Mb/s Ethernet ports:

500 UTP ports = \$37,500
 500 Fiber ports = \$260,000

<u>100 Mb/s cost/port</u>	<u>Twisted pair</u>	<u>Fiber</u>
Ethernet Hubs	\$100	\$330
End system cost	\$0	\$480 (100Base-T to fiber converter)
Cable cost (50 ft. cable)	<u>\$25 (Cat 5)</u>	<u>\$60</u>
Total	\$125	\$870

Assuming 40 100 Mb/s Ethernet ports:

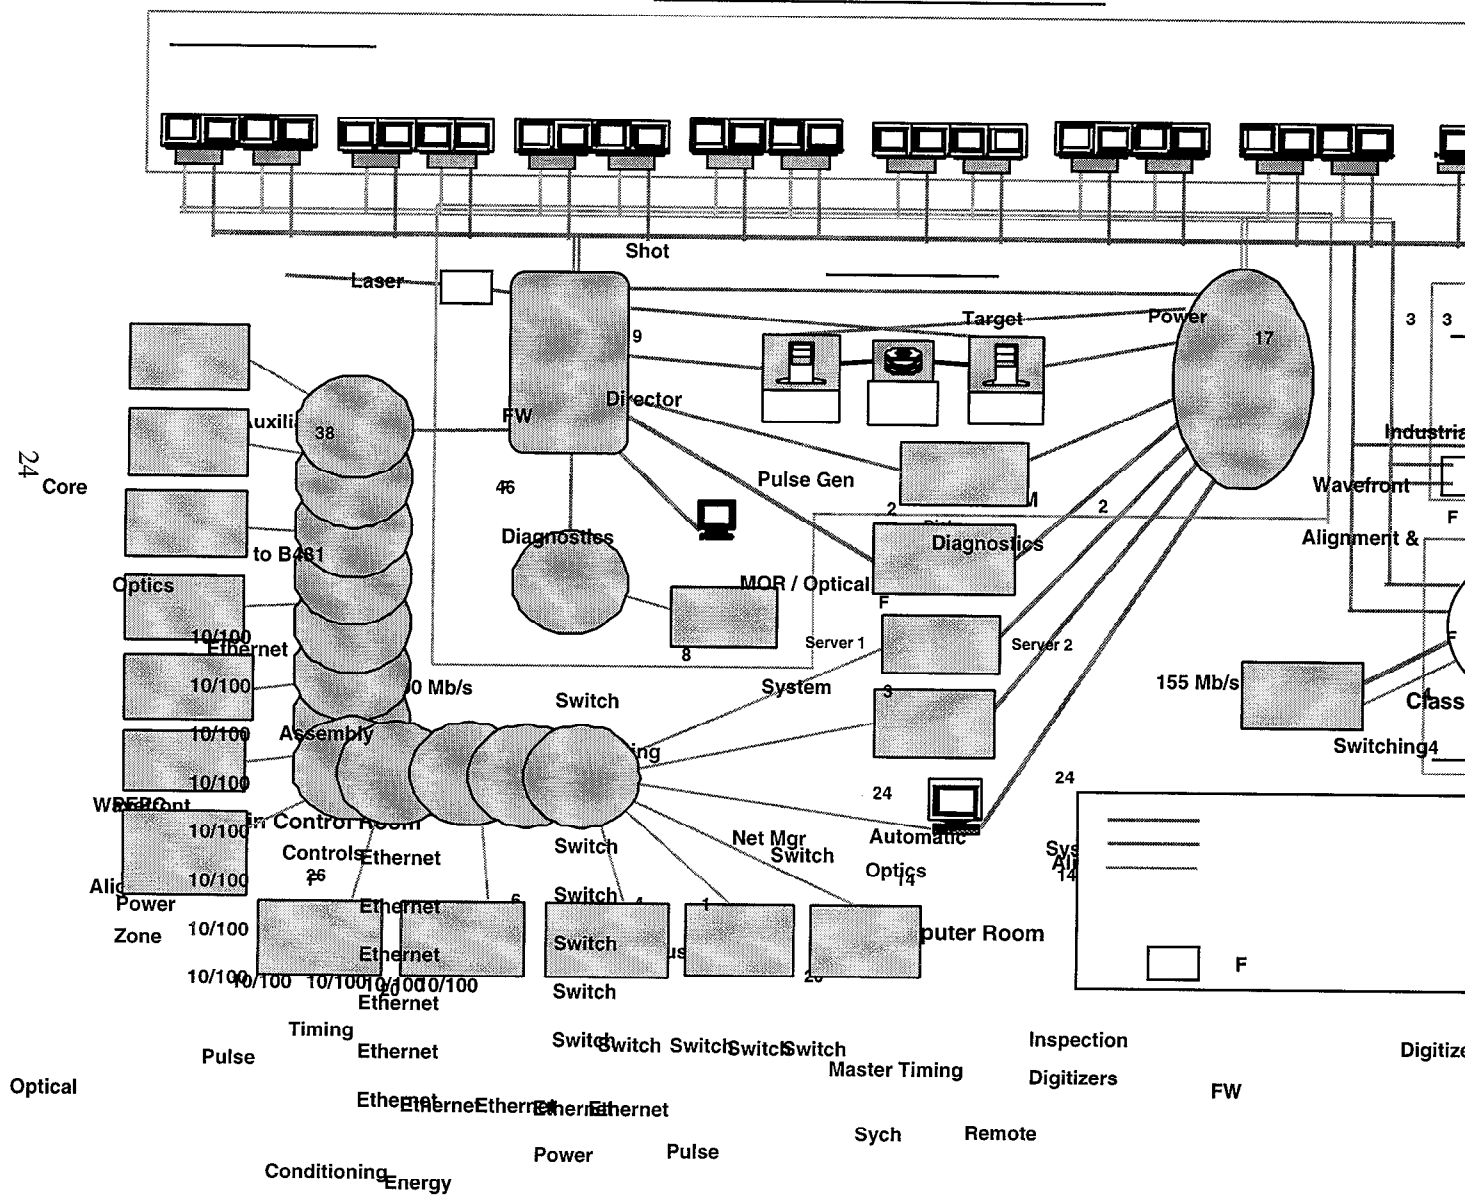
40 UTP ports = \$5,000
 40 Fiber ports = \$34,800

Extra cost to provide an all fiber solution to Ethernet end systems:

Fiber ports:	\$294,800
Twisted pair ports:	<u>\$42,500</u>
Difference	\$252,300

Network Switch MTBF Data

Vendor	System	Function	MTBF (yrs)	Notes
Cisco	Catalyst 5000	Chassis switch	7	system config.
	Catalyst 2820	Ethernet modular 2U switch	14.3	
	Catalyst 1900	24-port Ethernet Switch	14.3	
Cabletron	Smartswitch 9000	Chassis switch	49.4	calculated
			22.0	demonstrated
	Smartswitch 6000	Chassis switch, 5 modules	91.3	calculated
			11.8	demonstrated
	Smartswitch 2200	8-port FE Standalone switch	22.8	calculated demonstrated
	Smartswitch 2200	24-port Ether Standalone switch	19.9	
	Smartswitch 2200	24-port 10Base-FL Standalone switch	18.9	
Bay Networks	Smartswitch 2200	24-port 10Base-FL Standalone switch	22.8	
	Baystack 350T		75.8	
	Baystack 350FHD	24-port Ethernet switch	52.0	
	Accelar 1200	Switching Fabric	65.5	
		16-port 10/100TX	43.1	
		8-port 100FX	57.7	
Fore Systems	ASX-1000	ATM Switch (4 Fabrics)	3.2	
	OC-3 MM	4-port switch module	20.7	



Attachment: ICD Analysis Spreadsheets and Graphs